CHAPTER 5

Connecting to the Internet
(Growing into a Village)

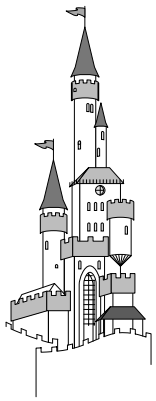

The Smiths picked a great location. So great, in fact, that others soon came to live nearby. Over time, their homestead turned into a group of farms and eventually grew into a small village. John Smith was elected the mayor of this village, and the village thrived. As the village grew, they built roads to other towns, villages, and cities around the area, hoping to encourage trade and communication. The roads also brought new dangers to protect against. Thieves out in the countryside threatened travelers. All kinds of people could ride into the village on the new roads, and some might not be trustworthy. This up-and-coming village had plenty of new trouble to watch out for.

Some residents thought to themselves, "Our homes are safe; those thieves are far away from here." Others thought, "It is easy to catch the bad guys, so we'll be protected. Besides, we have nothing they want." But John knew better. He talked to local officials and used the same thinking that had kept his home safe for all these years. John appointed a sheriff to help enforce the laws and allowed some of the residents to be deputies. The villagers built walls around some critical areas and added a strong vault for the bank, corrals for the horses, and barns and storehouses for food. The townspeople also all watched out for each other. They were neighbors and friends, who helped one another and kept an eye on unusual things. John even had the sheriff and the deputies ride the roads to check for trouble.

Types of Connections

You can connect your network or computer to the Internet in several ways. These involve plenty of differences, but also some important similarities. First, you must be running TCP/IP as your network protocol.

Why Should You Worry?

During the week of February 9, 2000, several of the biggest and best e-commerce sites (Buy.com, Amazon.com, Yahoo.com, and eBay.com, among others) were taken down in a Denial of Service (DoS) that was the first of its kind to hit so broadly. The DoS was generated in a distributed fashion, originating from literally hundreds of systems across the world and generating massive volumes of traffic. The sites were overwhelmed with the traffic, and eventually servers were unable to answer legitimate requests. Tracking the problem was difficult because the source of the traffic was computers that were unwitting accomplices. People like you and me owned those computers, as did large companies, universities, and many others. The original cracker planted "zombie" code on these boxes when they were unprotected and then later sent a simple command so the computers started sending network requests to the target. It was very effective.

Another example is a program that surfaced a while back called Back Orifice (BO) — supposedly a play on the name of Microsoft's Back Office. This program is a Trojan horse that allows the owner to do a wide variety of things on any system with this software on it. If I were running BO, I could attach to your system and open the CD tray, record your keystrokes, move your mouse, and more. That is pretty scary, but worse, then BO would publish your Internet address to a place where other hackers could find it and use your system too.

Let's say you are browsing the Internet and you get an e-mail message. The message appears to be from someone whose name you don't know, but the subject says "Here's that file we talked about." You're curious, so you open the mail and see that it says, "This one cracked me up, you should check it out." You figure it's some humor mail, probably from someone who knows you at the office, so you open the file. It takes you to a Web site and says "Loading . . . One Moment . . ."

At this point you might be perfectly safe, or you could be in big trouble. If this is a malicious hacker's attempt to compromise your system, they might well have succeeded. The mail was sent with a Trojan-horse program attached, and when you opened the file, it installed the program, possibly in addition to doing what was advertised or promised in the message. Now the

Why Should You Worry? *(continued)*

hacker can "visit" your system any time you are connected to the Internet. That's all the time if you're using DSL, cable, or ISDN, so the hacker has essentially unlimited use of your system. If you're able to control file access, you might stop some of these activities, but not all. A really tricky hacker can even send an e-mail message to your company as if it came from you, telling your boss you quit or that you want an outrageous raise.

You don't have to be running it on your system, but you do need to be running it at the point where you connect to the Internet. TCP/IP is the network protocol the Internet uses to operate. That means it is the language the Internet speaks. If you're wondering why this is important to security, think back to Chapter 3, Securing Your Computer, where we talked about protocol isolation. Not speaking the same language as everyone else increases your security. When you speak the same language, attackers already know some things about you and have a means of "talking" to your system. Attacks that can be carried out with only this knowledge are limited, but it is one less piece of information hackers must figure out before they can make an attack on your system.

Second, most people use the Internet in predictable and somewhat limited ways. By far, most people use e-mail, browse the Web, and maybe chat via IRC or instant-messaging services such as ICQ or AOL Instant Messenger. Coupled with the fact that most users are uneducated or lax about security, these predictable behaviors can be used to mount attacks against targeted networks or systems. An attacker who knows your behavior and what applications or protocols you're using most frequently can narrow down the number of things to try first in an attack. Your connection type is important because an attacker can only work when you are connected to the Internet. If your connection is always on and has a static IP address (one that doesn't change regularly), attackers have more hours per day to try to get in. You can see, then, that picking the right connection type and knowing its exposure is an important aspect of security. I'm not recommending that you move back to dial-up connections, but rather that you understand the

security issues involved with using the various connection types available today. Here are some of those issues:

- **Dial-up connection:** Using a standard phone line to dial in to an Internet service provider. This connection is not always present and often assigns Internet addresses (IP addresses) dynamically.
- **ISDN (integrated services digital network) connection:** An always-on connection that uses a special modem to connect at high speeds over dedicated lines. It can assign permanent addresses or dynamic ones, depending on the service provider.
- **DSL (digital subscriber line) connection:** Comes in two varieties: synchronous and asynchronous. (Their differences are beyond the scope of this book and not extremely relevant to security.) These are always-on connections that can assign addresses dynamically or statically, but usually statically.
- **Cable modem connection:** Runs through cables that used to carry only television signals but now carry network traffic too. Connections are often shared with other local cable users, but not always. IP addresses can be static or dynamic.
- **Satellite system connection:** Often configured to download from the satellite dish but upload across an attached modem and your phone line. Addresses can be static or dynamic and are not considered "always on."
- **WebTV/Internet appliance:** Generally connected through phone or cable connections. Often are just souped-up browsers with security equivalent to browsing the Web (discussed in later chapters).

What does it mean when we say IP addresses are static or dynamic? Static addresses are like your home address. Once you get an address, it stays with you until you move. An IP address is assigned by your Internet service provider (ISP) while you are getting your service through them. The ISP assigns dynamic addresses, too, but they have expiration dates and can change over time. The protocol for this is DHCP (dynamic host configuration protocol), which manages the assignment addresses from a pool of addresses used by the ISP. If you

want to investigate DHCP a bit more, you can find details about the full DHCP protocol in RFC 2131 at www.rfc-editor.org/rfc.html. (Some good information is also located at www.dhcp.org.) The security implications of static versus dynamic are often minimal. While it is true that a static address makes a computer easier to find on successive connection attempts, using DHCP doesn't make locating the target system that much more difficult. So in short, dynamic addresses are more secure, but only slightly so, and certainly not enough more secure that you don't need to use other security measures to protect your system.

Basic Internet Security

What can you do, then, to help secure your system when it's exposed to threats? Let's start with the basics. You need to take the following steps to secure your system when you're connected to the Internet. We've already covered many of these steps, so this is just a reminder. Remember that these are the foundations of good security, and if you do not follow these, all your other security measures will lose effectiveness.

- Secure your operating system to the best level it supports. To be truly securable, the OS must support user identities, security at the file-system level, and auditing of activities on the system.
- Don't run programs from unknown sources, including executing programs, scripts, or files containing macros.
- Use an antivirus program and be sure it scans your system regularly.
- Do not give out your password or logon information, and be careful with your personal information.
- Know your risk, and be aware of the value of the data on your system to yourself and others.
- Don't assume out-of-the-box security is enough.
- Turn on auditing if your operating system supports this option.

Advanced Internet Security

Now let's look at a few more advanced options for Internet security. These options are not required for most people; however, if you rated your risk as High, you should consider some or all of these options. (Again, we've talked about some of these in earlier chapters.)

- **Firewalls:** A firewall is some hardware—or a combination of hardware and software—that controls access to the traffic in and out of your network. Hmmm, sounds complex. Indeed firewalls can be very complex, but they can be simple too. Think of firewalls as the fences and gates that either allow traffic through or not. The typical home user doesn't need the power that most full-fledged firewalls offer. Instead, software packages called "personal firewalls" can serve the purpose for home users just fine. Generally speaking, these software packages should be capable of controlling outgoing and incoming traffic and setting "rules" concerning what traffic is okay and what isn't. They should also provide auditing or logging functions to let you determine if someone is trying to access your system without your permission. You can find more information about firewalls, as well as reviews and suggestions about which products are best for you and your situation, at www.firewallguide.com.

- **Proxy servers:** Different proxy servers will give you different functions, so I'll cover the basic concept first and then talk about some features you can find in these devices. Webster's dictionary[1] defines *proxy* as "authority or power to act for another," and that is exactly what happens here. A proxy server "acts on your behalf" on the Internet while your system sits behind the proxy, protected. All requests for Web pages, e-mail, chat, instant messaging, and such all are made from your systems to the proxy server. The proxy server then makes the request for your systems out to the Internet, without revealing your computer to the Internet. Attackers can't see your computer and potentially get access—they see only the proxy server. You

[1] *Merriam Webster's Collegiate Dictionary, Tenth Edition*. Springfield, MA: Merriam-Webster, Incorporated, 1993.

only have to secure the proxy, and the rest of your network can be protected behind it. If you have only one computer, don't bother with a proxy server; just protect the one computer. Additionally, some proxy servers offer *packet filtering*, which is the capability to block certain types of network traffic while allowing other traffic in. Some proxy servers act as complete firewalls, with incoming and outgoing filters, and some include auditing and logging of the traffic allowed and/or blocked.

- **Network address translation (NAT):** This very basic form of protection is essentially just hiding your address from the outside world. NAT acts like a proxy server for your address only. This is not very strong protection, but it is protection, and many of the newer Windows versions are shipping with this capability built in.
- **Audit log parsing:** Okay, you turn on your auditing so you can see what is happening on your system. That's good. But now you get a log full of events that are normal, and you have to sort through them to find the ones of interest. That's bad. This is a job for audit-log parsing tools. The name sounds complex, but they are usually easy tools to use. You tell them what events you want to see, and they search the logs and collect those events. The event logger in Windows NT and later versions can do limited filtering, but if you want the high-end stuff for systems at high risk, you can get parsing tools that can alert you to events in real time and can analyze events as they occur, trying to determine if the pattern is an attack or just normal activity. These advanced tools—called "intrusion detection programs"—might be a bit more than most homes and small businesses need, and they are usually costly. However, many of the personal firewall products available include these functions to some degree.
- **File encryption:** One of the oldest ways of protecting information is to encode or encrypt it. Romans used an encryption system to send messages between legions in big battles. They gave staffs of certain sizes to all commanders. Then they wound paper around a staff, wrote a message on the paper, and then unwound it. Only by having a staff of the correct diameter could someone rewind the paper and reconstruct the message. This

More About Encryption

You can use a program such as PGP[2] (which stands for Pretty Good Privacy) or Blowfish[3] to provide encryption for your e-mail. These programs use what is called public/private key encryption to accomplish their goals. This means you have one key that everyone in the world can know, and one key that only you know. When you encrypt a message with one key, it can be decrypted by using the other, and vice versa. Using this technology, you can protect messages from anyone but the intended recipient. Windows 2000 has an Encrypted File System (EFS) you can use to encrypt your files, or you can use third-party products to do the job if you are using other Windows-based systems. You can find some of these programs at www.tucows.com/system/fileencryption95.html.

It is important to know that no encryption is unbreakable. If you can encrypt a file, someone with enough computing power and time can decrypt it. The idea is to make the decryption so hard or time-consuming that it will do the person no good. For example, say you could somehow know who will win the 2015 World Series. You want to protect the information, so we'll encrypt it. At the time of this writing, 2015 is 13 years away Ñ roughly 177 million seconds (176,601,600, to be exact). If a person could guess once every second from now until 2015, that person would get 176,601,600 guesses at being right. We'll use a key to introduce randomness to the encryption, which allows us to control how strongly the data is encrypted. To protect our data, we want to make sure there are lots more choices than 177 million Ñ say, 100 times more Ñ so we choose a number between 0 and 20 billion (rounding up to make it even harder). Now, even by guessing once a second, a person has little chance of getting it right. Lucky for us, this simple example is a massive simplification of the real math done by people who do encryption, which means encryption can be both strong and safe.

One last thing about encryption: you might hear talk about encrypting and also about signing when referring to documents and files. Encrypting obscures the contents of the document or e-mail so that no one but the holder of the decryption key can read it. Signing, on the other hand, doesn't protect the document; it puts a block of encrypted text on the document as a signature. This block of text can be decrypted by your public key to show that it was indeed you that sent the document, much as a signature on a piece of paper or contract does.

[2] Freeware program developed by Philip Zimmermann

[3] Free program designed by Bruce Schneier

made the message reasonably secure in transit. Obviously, modern encryption is much more advanced, but it involves some of the same principles the Romans used. First you need a message or piece of data you want to protect. Second, you need a method for disassembling and reassembling the message reliably. Last, you need to ensure that all authorized parties know how to encrypt and decrypt properly and that they are the only ones who can. As a home user, the two places where you most likely would use encryption are for your e-mail and for your files.

- **Security Testing and Analysis Tools:** The last advanced option for Internet security is security testing and analysis tools. These tools are the same as or similar to the ones actual hackers use to access sites. I don't recommend this approach for novices because some of the tools can be complex; however, if you want (or need) to see how exposed you really are, try some of these tools on your systems. It can be an eye-opening experience. Some tools will deface Web pages, grant access to systems, load programs, let you literally control systems, or just leave a note saying you were there. These tools are the digital equivalent of a military training exercise. You'd better know how ready you are before you have to fight the battle, or you'll probably lose eventually. If you know where your weaknesses are, you can fix them, or at least protect yourself better. You can find a list of some security testing tools at www.insecure.org/tools.html.

Who Is Watching You?

With all of this talk about security, you might be wondering who is out there watching. What do they want with you? That question has many answers, and we'll explore them in the next few sections. But before we do, let me warn you that these sections touch on some areas that sound scary to most people. I have every intention of scaring you a bit with this information, but I don't want to scare you away. There are some rather unseemly characters out there in the world, and some of them are on the Internet. Locks and walls, doors, and maybe a dog can protect you at home. All I'm trying to do here is demonstrate that having protection on the Internet makes good sense too.

Let's say that now you are connecting to the Internet. You do so by dialing a phone or by using your cable or DSL connection. No one can possibly know you are there, right? Wrong. Let's hit the obvious ones first. Your ISP (Internet service provider) and the phone company or cable or DSL provider (if different than your ISP) all know you are connected. You haven't even done anything yet, and a few people already know. Of course, the Internet isn't really fun unless you do something, so next you hit the Web, answer some e-mail, and maybe start up your instant-messaging program. Now you've made some requests (called DNS requests) across the Internet to resolve names so you can get to those places. You've sent requests to Web sites, your online "buddies," and some other people through e-mail. What you might not know is that you've also sent information to the Web site owner and to advertisers through the banner ads that display on Web pages. Furthermore, your requests passed through probably dozens—possibly hundreds—of servers or routers along the way.

This is routine. There is nothing insidious or wrong about it; it is just the way the Internet works. But the point is that an attacker or someone who wishes to collect information about you (or anyone, for that matter) can "see" those requests and addresses and begin to get an idea of where you go and what you do on the Internet. With 10 billion billion addresses available on the Internet, you might think there are too many for anyone to "guess" yours, right? Wrong again. Those addresses are all between 0.0.0.0 and 255.255.255.255, and a knowledgeable programmer can use a computer to test each of those addresses at a rate of about several million a second. Also, some addresses are reserved and some are for special purposes, which reduces the number of required guesses. Eventually, someone will scan your address range and find you. It usually doesn't take more than a month of being online (it can be as short as a few hours or days) before someone "finds" that your address is live. Most people don't ever get beyond that, but some will try.

Wow—so people know you are out there. In fact, they probably know you are out there quite often. Who are these people? Most are businesses with legitimate reasons to know things: the people who carry the phone signal or run the devices that route Internet traffic, for example. The people that run the DNS servers to provide names of sites will know, if they choose to look. Such people are usually safe for two good

reasons. Because you are their customer, they already know a lot more about you than they can get on the Internet. They have billing addresses, phone numbers, and possibly credit card information (if you pay by that method). Also, they get literally thousands of DNS requests a second, which amounts to a huge amount of information every day. Even if they wanted to track it, doing so would take more time and money than the data is worth.

The advertisers and Web sites are a different matter, however. Banner ads you click on and Web sites you visit often glean e-mail addresses and browser information from you. UCE (unsolicited commercial e-mail or "spam") is big enough business to make lists of valid e-mail addresses valuable. Advertisers and marketing folks pay good money to know who is visiting the sites of products similar to theirs so they can try to cross-sell to you.

Yes, people are watching, and some of them are gathering information. But the last group are the ones to worry about most: the hackers and crackers and script kiddies. They use tools that are available on the Internet to watch addresses, try to break into systems, or attempt to disrupt things in general. They sometimes do it to be malicious, but sometimes it's just to see if they can.

Privacy Issues

With all of these people out there looking around, it's probably not surprising that your personal privacy is at risk. Advertisers and marketing people are always trying to gather more data so they can target their marketing to your tastes and introduce you to products that fit your lifestyle. (Of course, this is their picture of your tastes and lifestyle, based on snippets of information. Regardless of how good they are, they'll get some things wrong, and you'll be staring at advertisements that mean nothing to you.) But what can be worse are the people who collect your data to sell or who gather information about you that can be used in more harmful ways. For example, what if someone monitored the Web sites you visit and found you taking an interest in cancer information. They might report this to your insurance

company, who might raise your rates or drop your coverage for fear of having to pay for cancer treatment. This example is completely fabricated, but it could happen. There are many additional reasons for protecting this information, and, as they say, "Truth is stranger than fiction."

One example that actually occurs and often goes unknown for a long time is *identity theft*. If someone gets your Social Security number or taxpayer ID, they can get state identification as you in another state. With that, they can get credit cards, apartments, whatever—all in your name. They can request additional data about you by using this identification, and use that data to take out loans, buy cars, rent hotels, or travel. There is no real limit to what they can do, because to the rest of the world, they *are* you. As long as they stop using your identification and move on before you catch them, they can get away with this type of thing for a long time. People have had credit ratings ruined, houses foreclosed on, and incredible hassles from cases of identity theft, and this is only getting easier as more people use computers and have that data exposed online.

These examples show why you need to monitor your online privacy. The Platform for Privacy Preferences Project (P3P) from the W3C (World Wide Web Consortium) provides a set of rules that companies who build sites and software can use to help you control who gets what access to your information. You can learn more at the W3C Web site (www.w3.org/P3P/). Here's a quote from their site:

> The Platform for Privacy Preferences Project (P3P), developed by the World Wide Web Consortium, is emerging as an industry standard providing a simple, automated way for users to gain more control over the use of personal information on Web sites they visit. At its most basic level, P3P is a standardized set of multiple-choice questions, covering all the major aspects of a Web site's privacy policies. Taken together, they present a clear snapshot of how a site handles personal information about its users. P3P-enabled Web sites make this information available in a standard, machine-readable format. P3P-enabled browsers can "read" this snapshot automatically and compare it to the consumer's own set of privacy preferences. P3P enhances user control by putting privacy policies where users can find them, in a form users can understand, and, most importantly, enables users to act on what they see.

At this site you can also see products and tools available for taking advantage of P3P and what it does for you. This is a great step forward in Internet privacy and letting users take control of who gets access to information about themselves.

Internet Security Checklist

We have covered a lot of ground, so I have included a quick checklist here to help you assess how you are handling security for your Internet connection. This list should help determine how well you are covering the areas that need to be secured when you're connecting to the Internet.

- What type of connection do you have? Is it "always on?"
- Do you use an operating system that can be secured? Does it support user identities, security at the file-system level, and auditing of activities on the system?
- Do you run programs from unknown sources, including executing programs, scripts, or files containing macros?
- Do you use an antivirus program and make sure it scans your system regularly?
- Have you given out your password or logon information? Are you careful with your personal information?
- Do you know your risk and are you aware of the value of the data on your system to yourself and others?
- Do you assume out-of-the-box security is enough?
- Have you turned on auditing, if your operating system supports this option?
- Are you using a firewall, proxy, or network address translation (NAT)?
- Do you manually read audit logs or use a parser to do it?
- Do you protect sensitive information with encryption?
- Do you use any tools to analyze your own security? How often?
- Do you protect your online privacy?